

METHOD FOR SEQUENCING LANDING AIRCRAFTS

This invention relates to a method of sequencing vehicles. It has particular application for establishing the landing sequence of aircraft.

5

A phenomenon known as 'wake turbulence' is caused by wake vortices, which form whenever an aircraft wing is producing lift. The pressure differential between the top and bottom surfaces of the wing triggers the roll-up of the airflow aft of the wing resulting in swirling masses of air trailing downstream of the wing tips. The intensity or strength of the
10 vortices are primarily a function of the aircraft weight with the strongest vortices being produced by heavy aircraft.

Flying into the vortices can cause imbalance in following aircraft (possibly causing the following aircraft to crash) especially if the mass of the following aircraft is too small or the
15 intensity of the vortices is too great. As a result, a delay between two successive aircraft landings has to be maintained to avoid this potentially hazardous situation. This delay has to be extended proportionally to the mass ratio of the leading and following aircraft.

Assuming that there are three categories of aircraft ("heavy", "large" and "small") and that
20 the safe delay between them is an incremental function of their relative size, figure 1 shows a table summarising the delays (in time units, e.g. minutes) that must be maintained between successive landings. If all aircraft belonged to just one category then the delay would always be minimal. The delay would also be minimal if the arriving air traffic was grouped into three sets with all "small" aircraft landing first, followed by all
25 "large" aircraft and followed finally by all "heavy" aircraft. It is, of course, highly unlikely that timetable requirements would allow the organisation of air traffic into such a perfectly ordered sequence. In fact, aircraft belonging to all three categories follow each other at random and a problem facing air traffic controllers is choosing the aircraft which should be allowed to land next.

30

Two systems currently used by air traffic controllers to sequence incoming aircraft and to ensure landing aircraft are safely separated are Traffic Management Advisor (TMA) and Final Approach Spacing Tool (FAST) both developed by the National Aeronautics and Space Administration (NASA) Ames Research Centre, Moffett Field, California 94035,
35 USA.

Both these systems sequence incoming aircraft on a first come, first served (FCFS) basis where the first incoming aircraft to contact air traffic control (ATC) (with a request to land) on entering the terminal area (a term used to describe airspace in which air traffic control service is provided to aircraft arriving and departing an airfield) is allocated a landing slot first and placed at the start of the sequence. Subsequent, incoming aircraft are placed in the sequence in the order in which they enter the terminal area and contact ATC. Appropriate spacing is applied between sequenced aircraft to comply with safety constraints. It has been found, however, that sequencing aircraft on a FCFS basis leads to a less than optimal landing rate which leads to increased delays for arriving aircraft as they are forced to wait in the terminal area (usually in a waiting/holding stack) to be allocated a landing slot. This in turn leads to a reduction in quality of service provided by airlines and also to a increase in fuel consumption for the waiting aircraft.

15 According to a first aspect of the present invention there is provided a method of sequencing a plurality of candidate vehicles, wherein each candidate vehicle in said plurality of candidate vehicles is a candidate to be allocated the next place in a sequence, said method comprising the steps of:

- (i) receiving information pertaining to one of said candidate vehicles;
- 20 (ii) calculating a value to be attributed to said candidate vehicle on the basis of said received information and information received from the candidate vehicle most recently allocated a place in said sequence;
- (iii) repeating steps (i) and (ii) for each of said candidate vehicles;
- (iv) selecting one of said candidate vehicles based on said attributed values; and
- 25 (v) allocating said selected candidate vehicle the next place in said sequence.

Preferably the plurality of candidate vehicles comprises a plurality of candidate aircraft and the sequence is the landing sequence. By using information pertaining to candidate aircraft information from the aircraft most recently allocated a place in the sequence, a value can be calculated for each of the candidate aircraft and one of the candidate aircraft can be selected and allocated the next place in the sequence. The sequence of aircraft thus generated is more optimal than sequences otherwise generated, for example on a "first come, first served" basis.

Preferably, said received information is received from the candidate vehicle to which said received information pertains. In this way, it is more than likely that the received information will be up-to-date.

- 5 Preferably, said value is representative of the spacing that would have to be maintained between the candidate vehicle and the candidate vehicle most recently allocated a place in said sequence if said candidate vehicle were allocated the next place in the sequence. In this way, the average interval between successive vehicles is reduced.
- 10 Preferably, said value is representative of the delay that would be experienced by said candidate vehicle if said candidate vehicle were allocated the next place in the sequence. In this way, the average delay experienced by the candidate vehicles is reduced.

According to a second aspect of the present invention, there is provided a method of
15 operating a sequencing apparatus to sequence a plurality of candidate vehicles, wherein each candidate vehicle in said plurality of candidate vehicles is a candidate to be allocated the next place in a sequence, said method comprising the steps of:

- (i) receiving information pertaining to one of said candidate vehicles;
- (ii) calculating a value to be attributed to said candidate vehicle on the basis of said
20 received information and information received from the candidate vehicle most recently allocated a place in said sequence;
- (iii) repeating steps (i) and (ii) for each of said candidate vehicles;
- (iv) selecting one of said candidate vehicles based on said attributed values; and
- (v) allocating said selected candidate vehicle the next place in said sequence.

25

Preferably, said method further comprises the step of:

- (vi) sending details of the next place in said sequence to said selected candidate vehicle.

30 According to a third aspect of the present invention there is provided sequencing apparatus arranged in operation to sequence a plurality of candidate vehicles, wherein each candidate vehicle in said plurality of candidate vehicles is a candidate to be allocated the next place in a sequence, said data processing apparatus comprising:

- receiving means for receiving information pertaining to one of said candidate
35 vehicles;

calculating means for calculating a value to be attributed to said candidate vehicles on the basis of said received information and information received from the candidate vehicle most recently allocated a place in said sequence;

selecting means for selecting one of said candidate vehicles based on said
5 attributed values; and

allocating means for allocating said selected candidate vehicle the next place in said sequence.

According to a fourth aspect of the present invention there is provided sequencing
10 apparatus arranged in operation to sequence a plurality of candidate vehicles, wherein each candidate vehicle in said plurality of candidate vehicles is a candidate to be allocated the next place in a sequence, said data processing apparatus comprising:

a receiver arranged in operation to receive information pertaining to one of said candidate vehicles;

15 a calculator arranged in operation to calculate a value to be attributed to said candidate vehicles on the basis of said received information and information received from the candidate vehicle most recently allocated a place in said sequence;

a selector arranged in operation to select one of said candidate vehicles based on said attributed values; and

20 an allocator arranged in operation to allocate said selected candidate vehicle the next place in said sequence.

According to a fifth aspect of the present invention there is provided a digital data carrier carrying a program of instructions executable by processing apparatus to perform the
25 method steps as set out in the first aspect of the present invention.

Embodiments of the present invention will now be described, by way of example only, with reference to the accompanying drawings, wherein like reference numerals refer to like parts, and in which:

30

Figure 1 shows a table summarising the delays that must be maintained between successive landings of aircraft;

Figure 2 illustrates aircraft approaching a destination airfield;

Figure 3 illustrates a schematic view of the software used to implement an embodiment
35 the present invention;

Figure 4 is a flow diagram illustrating the first stages of an aircraft sequencing process;

Figure 5 is a flow diagram illustrating the remaining stages of an aircraft sequencing process;

Figure 6 is a flow diagram illustrating the calculation of a cost function in accordance with
5 an embodiment of the present invention;

Figure 7 is a flow diagram illustrating the computation of a landing time slot in accordance with an embodiment of the present invention.

Figure 8 is a table showing the results of sequencing aircraft on a "first come, first served" basis;

10 Figure 9 is a table showing the results of sequencing aircraft in accordance with an embodiment of the present invention;

Figure 10 is a graph showing a comparison in delays suffered by aircraft sequenced on a "first come, first served" basis and delays suffered by aircraft sequenced in accordance with an embodiment of the present invention.

15

In reference to figure 2, a plurality of aircraft 201 are shown approaching a destination airfield within a terminal area under the control of terminal area ATC 203. In order to request a landing time slot at the destination airfield, each of the aircraft 201 must contact terminal area ATC 203 upon entering the terminal area. The aircraft arrive in the terminal
20 area in an unpredictable fashion, i.e. in a random order.

A computer 205 within terminal area ATC 203 operates under the control of software executable to carry out an aircraft sequence selecting process. As will be understood by those skilled in the art, any or all of the software used to implement the invention can be
25 contained on various transmission and/or storage media such as floppy disk, CD-ROM or magnetic tape so that it can be loaded onto the computer or could be downloaded over a computer network using a suitable transmission medium.

Referring to figure 3, the software loaded onto computer 205 operates by attributing
30 and/or revising the priorities of entities ($E_1, E_2, E_3, \dots, E_n$) within a dynamic set 301. Associated with each entity (E_n) is a collection of real-time variables [$x(E_n), y(E_n)$]. The software further includes a scheduler 303 which operates in accordance with an optimisation algorithm in order to update the priority of the entities stored in the dynamic set 301 and move them to a static set 305. Each entity represents a single aircraft
35 arriving into the terminal area. Aircraft wait to be allocated a landing time slot in a

waiting/holding stack represented by the dynamic set 301. Using the optimisation algorithm, ATC (represented by the scheduler 303) decides the order of the landing sequence which is represented by the static set 305. Examples of the real time variables associated with each entity are the flight identification number of the aircraft, the size of the aircraft and the estimated time of arrival (ETA) of the aircraft at its destination.

Two further real time variables, I_n and D_n , are defined per entity for use in the algorithm by the scheduler. I_n is the interval to the aircraft represented by the latest entity in the static set 305 should the aircraft represented by entity E_n be allocated the next landing slot. (It will be remembered that this was described above in relation to figure 1.) D_n is the delay of the aircraft represented by entity E_n when compared with the aircraft's ETA should the aircraft represented by entity E_n be allocated the next landing slot.

The two variables, I_n and D_n , are combined into a cost function $f(I,D)$ which represents the associated 'cost' of allocating the next available landing time slot to the aircraft represented by the entity E_n . The relative weights of the two variables, I_n and D_n , in the cost function are adjustable and are defined as the value of two exponents, α and β . The cost function $f(I_n, D_n)$ is shown in full in equation [1] below:

$$f(I_n, D_n) = \frac{I_n^\alpha}{D_n^\beta} \quad [1]$$

The cost of selecting one entity from the dynamic set and transferring it to the static set (i.e. allocating the next available timeslot to an aircraft represented by entity E_n) is directly proportional to the interval, I_n raised to the power α and inversely proportional to the delay, D_n raised to the power β . A low interval and a high delay will decrease the cost of selecting a particular entity and hence decrease the cost of allocating a landing time slot to the represented aircraft. The longer an aircraft has already been waiting to be allocated a time slot, the more likely it becomes that it is allocated the next available time slot. However, all things being equal (i.e. all aircraft having similar delays), the aircraft with the shortest interval will be selected. This is best for maximising throughput of aircraft, reducing the chance of a long queue of waiting aircraft and therefore benefiting both the airfield and the aircraft.

Increasing α increases the weight of the interval I_n at the expense of the delay D_n . This typically results in minimal intervals between successive aircraft. On the other hand, increasing β increases the weight of the delay D_n at the expense of the interval I_n which typically results in reduced delays and hence reduced waiting times for incoming aircraft.

5

In preferred embodiments, the values of α and β are set to $\alpha=1.0$ and $\beta=2.0$. However, it is possible to modify the values of α and/or β to reflect changing priorities. Different aircraft may have different priorities due to, for example, an emergency situation on board the aircraft, the amount of fuel the aircraft is carrying, the total duration of the aircraft's journey, the nature of the cargo and/or the passengers onboard the aircraft etc. As a result, what is considered to be an 'acceptable delay' may vary accordingly.

10

Moreover, other variables and/or parameters could be added to equation [1] to account for other factors not included in the preferred embodiment, for example, the intrinsic priority of the aircraft, the current fuel consumption and/or fuel load of the aircraft, current atmospheric conditions, weather forecast etc. This would only modify the output variable returned by the cost function which is used as a decision basis by the scheduler.

15

In preferred embodiments, the decision as to which entity should be moved from the dynamic set to the static set and hence which aircraft should be allocated the next available landing time slot is made deterministically, that is, the entity with the lowest cost is moved.

20

Referring now to figure 4, on entering the terminal area, an approaching aircraft 201 contacts terminal area ATC 203 (step 401) via radio communication with a request for a landing time slot. This is assumed to take place anytime between ten and twenty minutes before the estimated time of arrival (ETA) of the aircraft at its destination. This initial contact message contains information such as a flight identification number of the aircraft, the size of the aircraft and the ETA of the aircraft. Upon receiving the contact message, terminal area ATC 203 acknowledges the message by sending a message back to the requesting aircraft 201 (step 403) which includes an order to wait in the waiting/holding stack. At the same time, an entity representing the requesting aircraft 201 is created by terminal area ATC 203 and added to the dynamic set 301 (step 405). This is achieved, for example, by inputting the relevant information onto computer 205 via a keyboard (or other such input device) attached to the computer 205. In other embodiments, the information

25

30

35

could be entered automatically into computer 205 via a datalink established between the requesting aircraft and terminal area ATC 203. Associated with that entity are the real time variables representing the information sent by the aircraft to terminal area ATC 203. The process described above in relation to figure 4 is repeated whenever an aircraft 201 enters the terminal area. Several aircraft 201 may enter the terminal area every minute contacting terminal area ATC 203 with a request for a landing time slot. This results in several entities being created and added to the dynamic set.

Referring to figure 5, the operation of the scheduler will now be described in further detail.

10 Firstly, a new session of the scheduler is initialised (step 501). A new session is begun for each landing time slot that is to be allocated by the scheduler. In preferred embodiments the scheduler is run once every minute although in other embodiments more or less sessions per minute may be more suitable.

15 The scheduler then extracts information (step 503) for the next entity representing an aircraft that has contacted terminal area ATC 203. The information extracted is that which the aircraft sent to terminal area ATC 203 in its initial contact message (figure 4, step 401). The scheduler then checks (step 505) whether or not the entity currently being processed has been waiting in the dynamic set for over a specified period of time, e.g.

20 thirty minutes. (It will be realised that this corresponds to an aircraft waiting in the waiting/holding stack for more than thirty minutes.) If this check yields a positive result then terminal area ATC 203 contacts the aircraft represented by this entity in order to re-direct it to another airfield (step 507) and the representative entity is removed from the dynamic set. If the check is negative then the scheduler continues to calculate the cost

25 function for this entity (step 509). The calculation of the cost function will be described in more detail below.

The scheduler then checks (step 511) whether or not the cost function just calculated is the lowest so far calculated in this session. If it is the lowest so far calculated then this

30 entity is temporarily classified as the best choice entity (step 513) until a time when the cost function of another entity is lower. Having calculated the cost function for the first entity in the current session, the scheduler then checks (step 515) whether or not cost functions for all the entities currently within the dynamic set have been calculated. If the result of this check is negative then steps 503 to 515 are repeated. If cost functions have

35 been calculated for all the entities currently within the dynamic set then the entity that

ends up classified as the best choice entity is moved from the dynamic set to the static set (step 517) and the scheduler computes (step 518) the next available landing time slot to allocate to the aircraft represented by the best choice entity. The computation of the landing time slot will be described in more detail below.

5

Having computed the landing time slot to be allocated to the aircraft represented by the best choice entity, the scheduler checks whether or not the delay associated with that aircraft (i.e. the difference between its allocated landing time slot and its ETA) is longer than a specified time period, e.g. sixty minutes. If the result of this check is positive then
10 terminal area ATC 203 contacts the aircraft in order to re-direct it to another airfield (step 521) after which time a new session of the scheduler is started. If the result of the check is negative then terminal area ATC 203 contacts the aircraft and informs it of its allocated landing time slot (step 523) at which time a new session of the scheduler is started.

15 With reference to figure 6, the calculation of the cost function (carried out in step 509) will now be explained in more detail. The scheduler first extracts information (step 601) from the last entity that was moved from the dynamic set to the static set. It will be realised that this entity represents the most recent aircraft to be allocated a landing time slot. The information extracted includes the size of the most recent aircraft and the landing time slot
20 allocated to it. Using this information and the size of the aircraft represented by the entity currently being processed (which it will be remembered was extracted in step 503), the scheduler then computes (step 603) what the interval (I) between these two aircraft would have to be if the aircraft represented by the entity currently being processed were allocated the next landing time slot. In the present embodiment, the intervals between
25 successive aircraft are those described above in relation to the table in figure 1, although otherwise defined intervals are also possible. The scheduler can then add this interval to the landing time slot allocated to the most recent aircraft to compute (step 605) a proposed landing time slot for the aircraft represented by the entity currently being processed. The scheduler can then compute the delay (D) (step 607) that the aircraft
30 represented by the entity currently being processed would suffer if allocated this landing time slot by comparing it with the aircraft's ETA. Finally the scheduler can use the interval I and delay D to compute the cost function (step 609) of the entity currently being processed.

With reference to figure 7, the computation of the landing time slot (carried out in step 518) will now be described in more detail. The scheduler first extracts information (step 701) from the last entity that was moved from the dynamic set to the static set. It will be realised that this entity represents the most recent aircraft to be allocated a landing time slot. The information extracted includes the size of the most recent aircraft and the landing time slot allocated to it. Using this information and the size of the aircraft represented by the best choice entity extracted by the scheduler in step 703, the scheduler then computes (step 705) what the interval (I) between these two aircraft has to be based on the intervals defined above in relation to the table in figure 1. Finally, the scheduler adds this interval to the landing time slot allocated to the most recent aircraft to compute (step 707) the landing time slot for the aircraft represented by the best choice entity.

It will be realised that in calculating the cst function for the best choice entity (in step 509), a proposed landing time slot for the aircraft represented by the best choice entity is calculated (in step 605). Hence in alternative embodiments, this information could be temporarily stored by the computer 205 and used by terminal area ATC 203 when it contacts the aircraft and informs it of its allocated landing time a lot (in step 523).

Figure 8 illustrates the landing sequence for the period 08:17 to 08:59 made on a "first come, first served" basis. Figure 9 illustrates the landing sequence for the same period and for an identical traffic pattern (same aircraft, same order of arrival) computed in accordance with the present invention.

The tables in both figures 8 and 9 are sorted by "Landing Time" which refers to the time the aircraft lands at its destination. "Flight ID" refers to the flight identification number of the aircraft, "Cat." refers to the size category of the aircraft, "ATC contact" refers to the time that the aircraft sends its initial contact message to terminal area ATC 203 on entering the terminal area, "ETA" refers to the aircraft's estimated time of arrival at its destination, "ATC allocate" refers to the time when terminal area ATC 203 contacts the aircraft with details of its allocated landing time slot and "Delay" refers to the difference in time between the aircraft's ETA and its actual landing time.

The shaded rows in the table 9 indicate aircraft that contacted terminal area ATC 203 earlier than some of the preceding aircraft but were allocated landing time slots later than

these predecessors. (This series of events can occur when the landing sequence is decided on a "first come, first served" basis but only when an aircraft that contacts terminal area ATC 203 has a later ETA than some of the following aircraft. This is indicated by the shaded rows in table 8.)

5

Referring to figure 8, thirty aircraft land in the forty-two minute period between 08:17 and 08:59. The average interval between them is one minute, twenty-five seconds and the average delay suffered by each aircraft is eighteen minutes, forty-four seconds. Referring to figure 9, thirty-seven aircraft land in the same forty-two minute period. The average
10 interval between them is now only one minute, ten seconds and the average delay suffered by each aircraft has fallen to fifteen minutes, sixteen seconds. This represents a 23.3% increase in capacity at the destination, a 18.5% reduction in the average delay suffered by arriving aircraft and a 17.4% reduction in the average interval between successive aircraft landings. This translates into a large improvement in quality of service
15 for the airlines operating the aircraft, including a substantial reduction in fuel consumption and an increase in revenue for airfields due to the increase in capacity.

The graph in figure 10 summarises the comparison. It is a plot of the delay suffered by aircraft against the time of day at which they land at their destination. By noon, nearly all
20 flights are delayed by at least thirty minutes and the situation continues to deteriorate since in the absence of any optimisation, the extra air traffic cannot be absorbed and the waiting/holding queue can only continue to grow. In contrast, the delay suffered by flights sequenced in accordance with the present invention remains fairly constant throughout the day. By the end of the day, three aircraft sequenced on a "first come, first served"
25 basis had to be re-routed to another destination because they suffered delays exceeding the maximum allowed delay (one hour in this case). The average delay suffered by aircraft was above thirty minutes compared with less than ten minutes for aircraft sequenced in accordance with the present invention.

30 Although in the above described embodiment the decision as to which entity should be moved from the dynamic set to the static set and hence which aircraft should be allocated the next available landing time slot is made deterministically, it is also possible to make the decision probabilistically on the basis of a function similar to:

$$C_x = \frac{f(I_x, D_x)}{\sum_{i=1}^N f(I_i, D_i)}$$

$$P_x = \frac{1 - C_x}{\sum_{i=1}^N 1 - C_i}$$

where N is the number of entities currently waiting in the dynamic set, C_x is the relative cost of selecting entity x and P_x is the probability that entity x is chosen.

- 5 Although the above embodiment was described in relation to the landing sequence of aircraft, it will be apparent that the present invention is just as applicable to the sequencing of any vehicles in a situation where those vehicles disturb the environment behind them as they proceed. One example of such a situation is ships/boats which leave a wake behind them.

10

The present invention successfully optimises sequences of vehicles. Test results suggest that sequencing aircraft about to land in accordance with the present invention leads to an increase in capacity at airfields (since aircraft can land more often) and an improvement to the quality of service provided by airlines operating those aircraft (since the delays suffered by aircraft is reduced). These two objectives were previously thought to be incompatible.

15